Using Deep Learning to Track Deforestation in the Amazon

Loren Amdahl-Culleton, Meredith Burkle, Miguel Camacho
CS 231N: CNNs for Visual Recognition, Spring 2017

Motivation

According to National Geographic, almost one fifth of the Amazon rainforest has been cut down in the last 40 years. However, comprehensive understanding of the changes is lacking as traditional models provide *inaccurate estimates of the extent of the deforestation*.

Existing methods fall short for various reasons including an inability to *differentiate between different types of deforestation* and the inability to make use of *improved satellite image resolution*.

In order to improve forest image classification, we will apply state-of-the-art *multi-class, multi-label CNN* frameworks to novel Planet Labs satellite images. By building a robust and accurate classifier, we hope to allow a better understanding of the changes in the Amazon rainforest to ultimately inform public opinion and public policy.

Data

To train our model we used the data available through the Amazon Kaggle challenge, which contains more than 40K labeled satellite images, or "chips," which describe the land use, land cover, atmospheric cover and weather patterns depicted in the image.

Training Set	~32K
Validation Set	~4K
Development Set	~4K
(Hidden) Test Set	~40K

In order to augment our data for a more robust classifier we distorted our training images with 50% probability in various ways including:

- Randomly mirroring images horizontally
- Randomly scaling images by 10%
- Randomly multiplying pixel values by 5% (manipulating brightness)
- Randomly cropping images by 5%

Model

Baseline:

- Custom 2 Layer CNN with Max Pooling and ReLU activation to get sanity checks and baseline results.
- The model trains using sigmoid classification and BCE Loss.

Inception-SS (Softmax-Sigmoid):

- Used Inception net weights to extract image features
- Used multi-label sigmoid classification over land-type labels, single-label softmax classification over weather labels with cross-entropy loss

Inception-LSTM:

- Extract image features using Inception net
- Feed image feature vector into LSTM architecture that predicts a series of labels ("captions") for each image

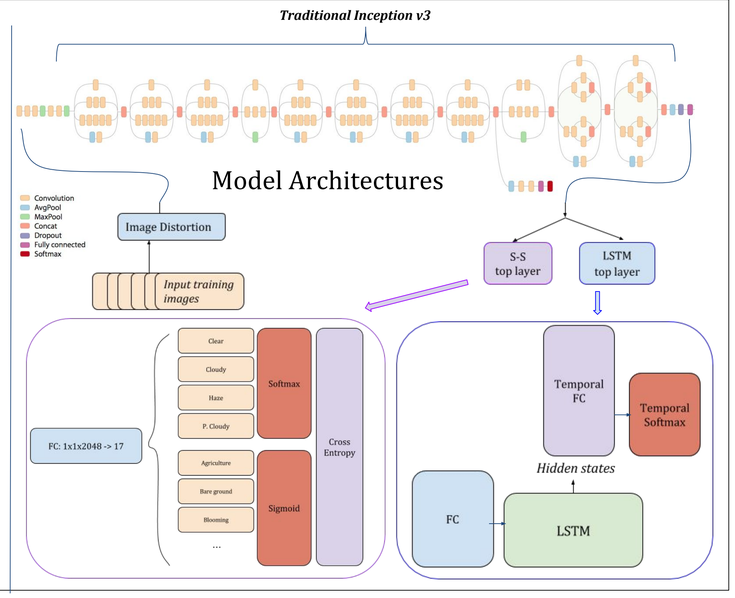

Performance

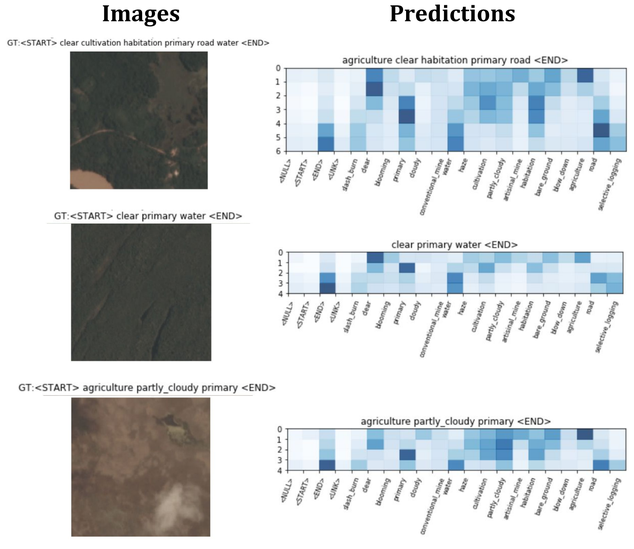

The plots to the left show the **probability distributions** at each time step. The darker the square, the higher the probability.

Starting from the top, we see the first predicted label, continuing until the last row, where the model predicts the '<END>' token.

The **top** image and prediction pair display a more difficult label set, as we can see the model is not very confident in its predictions.

The **middle** image/prediction pair display a relatively common label set, and we can see that the model is confident in its predictions.

The **last** pair showcases the model's ability to predict weather patterns. Despite the image showing similar characteristics to a 'haze' or 'cloudy' image, it correctly (if hesitantly) guesses 'partly cloudy.'

Evaluation Metrics

Model Accuracies

	Train (Accuracy/F2)	Val (Accuracy/F2)
Baseline	0.90/0.64	0.91/0.65
Inception-SS	(still training as of 6/5 11:59pm)	0.93/0.74
Inception-LSTM	0.96/**0.89**	0.95/**0.87**

The top 5 teams on the Kaggle leaderboard have attained F2 scores of above 0.932.

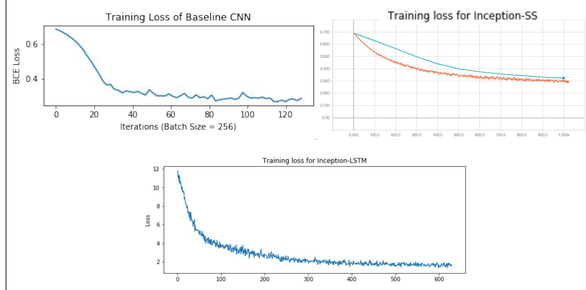

Conclusions

- We proposed a transfer learning scheme based on Inception net that was modified to predict one weather label and one or more land labels and achieved a ***best one-to-one accuracy of 95%*** and a ***best F2 accuracy of 87%***
- Our current architecture decisions have improved loss performance as well as F2 accuracy, but we are still working on re-training deeper pre-trained layers to achieve F2 scores competitive with top Kaggle Challenge submissions
- We are currently working on incorporating TIFF files to our training data and making use of the NIR data channel in initial preprocessing

Acknowledgements:

- Entire CS231N staff!
- Planet Labs
- Kaggle